



INFORMATION TECHNOLOGIES  
& SERVICES DIVISION

# Scientific Cluster Support Project

2003-2004 Activities, Challenges, and Results

Gary Jung  
SCS Project Manager

January 7, 2005

---

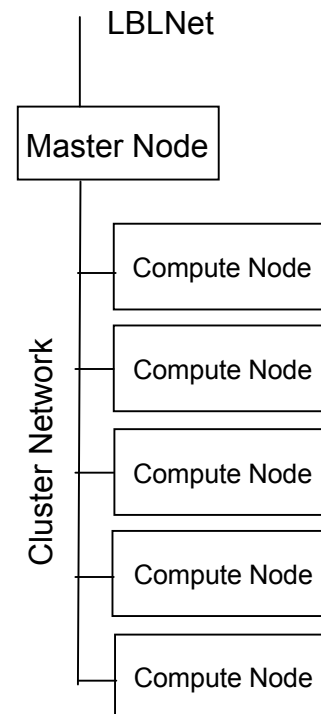


- **Why is scientific computing so important to our researchers?**
  - **Traditional methods**
    - Theoretical approach
    - Experimental approach
  - **Computational approach is now recognized as important tool in scientific research**
    - Data analysis
    - Large scale simulation and modeling of physical or biological processes

# A Brief History of Computing at Berkeley Lab

- **The 1970's and early 1980's – Central computing**
  - CDC 6000 and 7600 Supercomputers
- **The 1980's – Minicomputers**
  - Digital Equipment Corp VAX and 8600 series systems
  - Interactive timesharing computing
- **The 1990's – Distributed networked computing**
  - Computing at the desktop
  - Institutional central computing fades away
  - The “Gap”
- **2000 - Linux cluster computing starts to emerge at Berkeley Lab**

- **Commodity Off The Shelf (COTS) parts**
- **Open source software (Linux)**
- **Single master/multiple slave(compute) node architecture**
  - External view of the cluster is as a single unit for managing, configuration, communication
  - Organized dedicated network communication among nodes
- **Similar or identical software running on each node**
- **Job scheduler**
- **Parallel programming software - Message Passing Interface (MPI)**



# Scientific Cluster Support Project Initiated

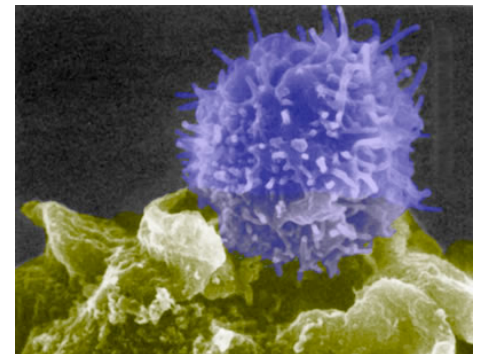
- **2002 - MRC Working Group recommends that ITSD provide support for Linux Clusters.**
- **December 2002 - SCS Program approved**
  - \$1.3M Four-year program started January 2003
  - Ten strategic science projects are selected
  - Projects purchase their own Linux clusters
  - ITSD provides consulting and support
- **Strategy**
  - Use proven technical approaches that enable us to provide production capability
  - Adopt standards to facilitate scaling support to several clusters
- **Goals**
  - More effective science
  - Enable our scientists to use and take advantage of computing
  - HPC that works. Avoid lost time and expensive mistakes

# Participating Science Projects

Chemical Sciences	PI: William Miller	Semiclassical Molecular Reaction Dynamics: Methodological Development and Application to Complex Systems	40 Intel Xeon processors
Chemical Sciences	PI: Martin Head-Gordon	Parallel electronic structure theory	42 AMD Opteron processors
Chemical Sciences	PI: William Lester	Quantum Monte Carlo for electronic structure	46 AMD Athlon processors
Materials Sciences	PI: Arup Chakraborty	Signaling and Mechanical Responses Due to Biomolecular Binding	96 AMD Athlon processors
Material Sciences	PI: Steve Louie Marvin Cohen	Molecular Foundry	72 AMD Opteron processors
Physical Bioscience	PI/Contact: Kim/Adams/ Brenner/Holbrook	Structural Genomics of a Minimal Genome Computational Structural & Functional Genomics A Structural Classification of RNA Nudix DNA Repair Enzymes from <i>Deinococcus radiodurans</i>	60 Intel Xeon processors
Environmental Energy Technologies	PI: Gadgil/Brown	Airflow and Pollutant Transport in Buildings Regional Air Quality Modeling Combustion Modeling	24 AMD Athlon processors
Earth Sciences	PI: Hoversten/Majer	Geophysical Subsurface Imaging	50 Intel Xeon processors
Life Sciences	PI: Michael Eisen	Computational Analysis of cis-Regulatory Content of Animal Genomes	40 Intel Xeon processors
Life Sciences	PI: Cooper/Tainer	Protein Crystallography and SAXS data Analysis for Sibyls/SBDR	20 Intel Xeon processors
Nuclear Sciences	PI: I-Yang Lee	Gretina Detector - Signal deposition and event reconstruction	16 AMD Opteron processors

- **Scheduling**
  - Funding availability
  - Variance in customer readiness
- **Security**
  - Export control
  - One-time password tokens
  - Firewall
- **Software**
  - Licensing LBNL developed software
  - Red Hat Enterprise Linux

- 14 clusters in production
  - 10 SCS funded, 3 fully recharged, 1 ITSD test cluster
  - 698 processors online
- Warewulf cluster software
  - Standard SCS cluster distribution
  - University of Kentucky KASY0 supercomputer
- ITSD at Supercomputing 2003
- Enabling science
  - Chakraborty T-cell discovery - Oct 2003
  - Lester INCITE work on Photosynthesis - Nov 2004



- **Driving down costs**
  - Standardization of architecture and toolset
  - Outsourcing of various pieces
  - Develop lower cost staff
  - Competitive bid procurement
    - About 10% savings
  - Benchmarking costs
    - Comparison to postdocs
    - Comparison to other Labs

- Initial funding was key to get started
- Prominent scientists were our customers
- Talented, motivated staff
  - Creative, but focused on production use
  - Development of technical depth
- Adherence to standards
- Supportive Steering Committee
- Positive feedback

- **Larger systems**
  - Scalability issues - e.g. parallel filesystems
  - Moving up the technology curve - Infiniband, PCI Express
  - Assessing integration risks
- **Increasing cluster utilization**
- **Harder problems to debug**
- **Charting path forward**

- **Upcoming projects**
  - Earth Sciences 256 processor cluster - Spring 2005
  - Molecular Foundry 256 processor cluster - Dec 2005
  - Gretina 750 processor cluster 2007
- **Follow-on to SCS**
  - SCS approach vs. large institutional cluster
  - Grids

**PI: Arup Chakraborty**  
**Materials Sciences Division**

96 AMD 2200+ MP processors  
48 GB aggregate memory  
1 TB disk storage  
Fast Ethernet interconnect  
345 Gflop/s (theoretical peak)



**PI: Steve Louie and Marvin Cohen**  
**MSD Molecular Foundry**

72 AMD Opteron 2.0 Ghz 64-bit processors  
72 GB aggregate memory  
2 TB disk storage  
Myrinet interconnect  
288 Gflop/s (theoretical peak)

